

IPv6 Operations - IETF 110
2021-03-09 17:00 UTC

Chairs: Fred Baker, Ron Bonica

AD: Warren Kumari

A breif discussion on the use of IPv6 in the Internet

Fred and Ron...

Slides: [Chairs thoughts on IPv6 Deployment](#)

- Fred mentions Class-E space discussions, active in NANOG "its a band-aid"
- Mentions the divergence in measurements: Akamai, Google measure to the CDN. APNIC measures e2e under transit internet-wide.
- Argues "it costs less money, TCO" (for a new deployment which would have to buy IPv4 addresses)
- Overlay/Underlay, is it possible that the model is causing a "are we doing something wrong" moment

Fred: No request for discussion at this time, because it will come up in later talks in this session. But I'd like people to think about whether we, as IETF, are doing something collectively wrong?

Processing of the Hop-by-Hop Options Header

2021-01-21, draft-peng-v6ops-hbh, Gyan Mishra

Slides: [Processing of the Hop-by-Hop Options Header](#)

Ron Bonica: Very important. Because of the WG Charter. Parts of it falls in the 6MAN. Parts of it falls outside of our Charter.

Gyan: Understood. We do have draft on HBH forwarding. It's listed on the slide. That's in 6man. This draft is just the problem statement. The solution draft is in 6man.

Ron: This draft needs to be devoid of solution.

Bob: I have draft in 6man that proposes different solution. We need to make HBH processing simpler. The notion of creating another option for HBH will just make things more complicated and will not solve this problem. I support v6ops doing problem statement draft, but it needs to be devoid of solution.

Gyan: Understood.

Jen Linkova: I think that what sorts of things that can be processed is a property of the router. It will change over time because routers will be able to process more. How does this draft address that?

Gyan: It's evolving. If more options come up that need to be in slow path, they would remain in slow path. It's 2 separate buckets and that's something we would have to consider. It's 2 buckets. One way to look at it - if you have existing options, there would be no change. New things that require line rate processing in the fast path would go in the fast lane. Existing control plane capabilities stay in the slow lane.

Fred: Please ditch "fast path" and "slow path". It's about control plane and forwarding plane. If you use those terms, the discussion becomes clearer.

Gyan: I can see that. OK. I think I agree that having this draft focus on the problem statement makes sense.

Fred: Looking at charter - the charter allows us to develop problem statement and send them to others. And also to comment on operational aspects of solutions. I'd like to see the draft be reworked and literally be a problem statement. Then we can discuss adoption.

Gyan: Understood.

Ron: I agree we have a real problem. We have so many proposals that need more bits. HBH extension would be natural solution. But we have so many that it's being shoe-horned in. To keep this draft in scope of v6ops, we do need problem statement devoid of solution.

Gyan: Understood. We'll do that and take to mailing list for review.

Fred: Thanks. Shuping has re-entered queue and can now be heard.

Shuping: We share same goal, as Bob said. We want to make HBH useful for operators. Currently we have 2 drafts, one here and one in 6man. We do have this draft here for the problem statement. We will remove the solution part. We do also think it would be useful to have requirements for solution. We would like to get feedback from WG.

Warren Kumari (no hat): If solution is removed, it seems like all you're left with is "it sure would be nice if HBH worked". There is so much vendor-specific functions, that it's not clear how this will change anything. It just seems to be "I want a pony".

Gyan: With changes to draft, we'll focus on really defining the problem statement. I think there is something here we'd like to address. We're open to all solutions that can solve the problem. We'll focus on defining the problem in this draft.

Bob: I agree with a lot that Warren said. 2 thoughts... If we can get something that works we need to be careful about what we allow to be defined for new HBH options. I could easily imagine defining new things that don't need to go in the forwarding path (fast path). We also need some indication from operators that this is something they think they'll deploy and use. We need collaboration with the people who build routers and chips and with people who operate networks.

Ron: To argue a bit with Warren... I think ASICs are becoming more powerful and may be capable of doing more.

Shuping: Just a clarification to Warren. It's not about the operators, but it is about the implementation in the old times when the hardware is not capable, just as Ron commented. The problem is that when the hardware becomes capable the implementation still stays there with no change. That caused the trouble. There is a need to guide the new implementation, but currently there is still no such specification. We need something to fill the gap. That is the purpose of the draft in the v6ops, while 6man works on the solution.

Gyan: The big problem is chip manufacturers. But there are 2 problems. Some ASICs that are not as capable and other chips that don't have functions.

IPv6 Deployment Status

2021-02-22, draft-vf-v6ops-ipv6-deployment, Paolo Volpato

Slides: [IPv6 Deployment Status](#)

Fred Baker: A few minutes for comments.

Ron (no hat): Questions... An increase in IPv4 price and emergence of v6-only islands would motivate adoption. Have we seen emergence of v6-only islands?

Fred: The answers to both questions are Yes. We've seen increase in price of v4 addresses (12 to 20 dollars) and emergence of v6-only islands.

TCPM RTO-dependent flow label generation discussion

Alexander Azimov

Slides: [TCP Socket Hash & Flow Label](#)

Gyan Mishra: Good overview. No flow label in IPv4, we need to take advantage of this when deploying IPv6.

Alexander: Yes, we did flow label balancing on different devices. On some, flow labels were not working properly and vendors have provided us with updates. We will be testing with these updates.

Michael Tuxen: Prefer to disable in the server side. Since packet has no states. Not talking about Anycast.

Alexander: They're not normally in anycast. If you replace the firewall in your gateway, everything will be fine. On client side it will be fine. On server side it may time out.

Michael: I'm not talking about anycast. The packets might just not hit the firewall and the packets will die. Firewall on the path. May not hit the FW anymore.

Alexander: What happen in the scenario?

Michael: Temporary problem. It changes routing.

Alexander: Changing path if you have multiple FWs. Flow label is already changing.

Luigi Iannone: Interesting presentation. What happens ... Is there time to do an experiment on this?

Alexander: Yes, we are currently running some experiments. For content provider, problem is not normally on ingress but on egress. You are right that if there is packet loss on forward path from service provider this will limit ability of packet to jump to other path. The safe mode is an easy fix. There is negotiation between client and server. But safe mode should be deployed.

Luigi Iannone: How about MPTCP? Same principle? If complex we go to the list.

Alexander: It's not an easy question. Agreed to take this question to the list.

Gyan: Since you've had erratic results with stateful marking of flow label, does it make sense from the operator perspective to unmark the packets so flow label is not set. It's untrusted and stateful isn't used?

Alexander: I didn't fully follow you. We never remark flow label at any point in the network. It's just about the function of the ingress device. I'm not aware of other stateful uses of the flow label.

Gyan: But on ? rack of the ingress point, could you set it so field isn't marked so you don't have erratic behavior?

Alexander: ...

Fred Baker: We need to move on.

IPv6 Addressing Considerations

2021-02-21, draft-gont-v6ops-ipv6-addressing-considerations

Slides: [IPv6 Addressing Considerations](#)

Jen Linkova: Very interesting. Different profile for different ID. Cover a lot of different things. Agree with API. Not sure about how to communicate those properties.

Fernando: One of the options (but this just scratches the surface)... At the time temporary addresses were specified, there was something that allowed a socket to indicate it wanted a temporary address. I don't know if this is the solution, and this is just me thinking aloud. There might need to be some more IPv6-specific specification done.

Fred: Are you suggesting we change the TAPS charter or something else?

Fernando: First we need to agree on the problem. We could have a doc in v6ops that talks about the properties and what needs to be exposed. But if we agree on the problem, we need discuss with TAPS what would be the best way to proceed so they would have what they could to proceed with protocol development. TAPS doesn't know the IPv6 details so they couldn't do the whole thing.

Fred: We will explore how to go from there. We will need to figure out the best way to cooperate with them. I would expect from ART area. Maybe we could provide requirements.

Fernando: People who uses it don't have to know about the details. This doc is problem statement.

Luigi Iannone: Good document and presentation. Thanks. I'm aware of another problem and would like to talk to you offline because I think it's inter-related.

Bob Hinden: Worthy pursuing. Large doc now trying to convey a lot of different things, try to pull out stuff and separate.

Fred: OK. Before we go... In last several meetings, Ron and I have asked the WG to talk about a different draft per week. Is this useful? I'll make this a hand-raise thing. // 13 of 47 participants raised hand to say they found it useful. One raised hand to say they did not find it useful.